

Constrained Thompson Sampling for Real-Time Electricity Pricing With Grid Reliability Constraints

Nathaniel Tucker¹, *Graduate Student Member, IEEE*, Ahmadreza Moradipari,
and Mahnoosh Alizadeh¹, *Member, IEEE*

Abstract—We consider the problem of an electricity aggregator attempting to learn customers’ electricity usage models while implementing a load shaping program by means of broadcasting dispatch signals in real-time. We adopt a multi-armed bandit problem formulation to account for the stochastic and unknown nature of customers’ responses to dispatch signals. We propose a constrained Thompson sampling heuristic, Con-TS-RTP, as a solution to the load shaping problem of the electricity aggregator attempting to influence customers’ usage to match various desired demand profiles (i.e., to reduce demand at peak hours, integrate more intermittent renewable generation, track a desired daily load profile, etc). The proposed Con-TS-RTP heuristic accounts for day-varying target load profiles (i.e., multiple target load profiles reflecting renewable forecasts and desired demand patterns) and takes into account the operational constraints of a distribution system to ensure that customers receive adequate service and to avoid potential grid failures. We provide a discussion on the regret bounds for our algorithm as well as a discussion on the operational reliability of the distribution system’s constraints being upheld throughout the learning process.

Index Terms—Constrained optimization, distribution network, multi-armed bandit, real-time pricing, demand response, Thompson sampling.

I. INTRODUCTION

IN ORDER to integrate the increasing volume of intermittent renewable generation in modern power grids, aggregators are exploring various methods to manipulate both residential and commercial loads in real-time. As a result, various demand response (DR) frameworks are gaining popularity because of their ability to shape electricity demand by broadcasting time-varying signals to customers; however, most aggregators have not implemented complex DR programs beyond peak shaving and emergency load reduction initiatives. One reason for this is the customers’ unknown and time-varying responses to dispatch signals, which can lead to

economic uncertainty for the aggregator and reliability concerns for the grid [1]. The aggregator could explicitly request price sensitivity information from its customers; however, this two-way negotiation has a large communication overhead and most customers cannot readily characterize their price sensitivities, and even if they could, they might not be willing to share this private information. As such, aggregators prefer the 1-way *passive* approach because it does not require any real-time feedback from the customer and it does not require new communication infrastructure for reporting preferences (e.g., a Web portal, phone application, etc.). With this in mind, future load shaping initiatives for renewable integration (i.e., more complex objectives than peak shaving) need to be able to *passively* learn customers’ response to dispatch signals only from historical data of past interactions [2].

Recently, much work has been done for aggregators attempting to learn customers’ price responses whilst implementing peak shaving DR programs. The authors of [3] present a data-driven strategy to estimate customers’ demands and develop prices for DR. In [4], the authors use linear regression models to derive estimations of customers’ responses to DR signals. Similarly, [5] develops a joint online learning and pricing algorithm based on linear regression. In [6], the authors present a contract-based DR strategy to learn customer behavior while broadcasting DR signals. The authors of [7] present an online learning approach based on piecewise linear stochastic approximation for an aggregator to sequentially adjust its DR prices based on the behavior of the customers in the past. In [8], the authors develop a risk-averse learning approach for aggregators operating DR programs. In [9], a learning algorithm for customers’ utility functions is developed and it is assumed that the aggregator acts within a two-stage (day-ahead and real-time) electricity market. Additionally, the authors of [10] present a learning framework for forecasting individual loads and DR capabilities and find that users with more variable consumption patterns are more effective DR participants. Using a similar framework as in this work, a multi-armed bandit (MAB) formulation is used in [11], [12] to determine which customers to target in DR programs.

In addition to learning how customers respond to DR signals, an aggregator must also consider power system constraints to ensure reliable operation (e.g., nodal voltage, transformer capacities, and line flow limits). In real distribution systems, it is critical that these constraints are satisfied at every time step to ensure customers receive adequate service and to avoid potential grid failures even without sufficient knowledge

Manuscript received August 20, 2019; revised January 3, 2020 and April 23, 2020; accepted May 19, 2020. Date of publication June 25, 2020; date of current version October 21, 2020. This work was supported in part by NSF under Grant 1847096 and Grant 1737565, and in part by UCOP under Grant LFR-18-548175. Paper no. TSG-01223-2019. (Nathaniel Tucker and Ahmadreza Moradipari contributed equally to this work.) (Corresponding author: Nathaniel Tucker.)

The authors are with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: nathaniel_tucker@ucsb.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSG.2020.3004770

about how customers respond to price signals (i.e., in early learning stages) [13], [14]. One paper similar to ours that considers these realistic constraints, [15], presents a least-square estimator approach to learn customer sensitivities and implements DR in a distribution network. The authors of [15] show that their least-square algorithm's parameter estimation error converges to zero over time, thus the algorithm's regret is sublinear while also accounting for the distribution network's constraints.

Similar to the aforementioned papers, the work presented in this manuscript considers the problem of an aggregator *passively* learning the customers' price sensitivities while running a load shaping program. However, our approach permits more complex load shaping objectives (e.g., tracking a daily target load profile) and varies in terms of both load modeling and learning approach from all the above papers. Specifically, we present a multi-armed bandit (MAB) heuristic akin to Thompson sampling (TS) to tackle the trade-off between exploration of untested price signals and exploitation of well-performing price signals while ensuring grid reliability. It is important to note that the standard TS heuristic cannot guarantee that grid reliability constraints are upheld during the learning process. As such, we present two *modified* versions of TS while retaining the fundamental principles TS is based on. Furthermore, we provide discussion on how the constraints are upheld (i.e., operational reliability) for the modified heuristics, discussion on the performance of the heuristics compared to a clairvoyant solution, and simulation results highlighting the strengths of the method.

In our work, we make use of a load clustering technique in order to *exploit the known physical structure* of the problem and make use of our prior knowledge of how flexible electric appliances behave to lower the problem dimensionality. We note that grouping (clustering) loads for dimensional reduction is common in DR literature [16]. Some pertinent examples include [17] where the authors aggregate heterogeneous thermostatically controlled loads (TCLs) using an LTI "bin" model, [18] where the authors group EVs into "classes" depending on their charging availability, [19] where the authors present a load profile clustering method for load data classification based on information entropy, piecewise aggregate approximation, and spectral clustering, [20] where the authors present aggregate models for classes of TCLs that include statistical information of the population, systematically deal with heterogeneity, and account for a second-order effects, [21] where the authors propose a clustering technique for determining natural segmentation of customers and identification of temporal consumption patterns in the smart grid domain, and [22] where the authors develop cohorts, or groups of consumers with similar consumption patterns, from correlations between daily loads.

The **main contributions** of this work are as follows:

- We use the multi-armed bandit (MAB) framework to model the stochastic and unknown nature of customers' daily aggregate response to electricity prices.
- We make use of an appliance clustering methodology to provide a mesoscopic model of the price responsive

demand of a large population of flexible appliances and reduce the dimensionality of the learning problem.

- Our learning framework can account for daily variabilities and realistic grid reliability constraints that are critical for daily operation in spite of uncertainty about customers' price response.
- We present two modified heuristics based on Thompson sampling (TS) as solutions to the constrained learning and pricing problem.
- We provide a performance guarantee in the form of a regret bound and discussion on the reliability guarantees of the approach as well as a distribution system case study demonstrating the efficacy of the approach.

The remainder of the paper is organized as follows: Section II presents the aggregator's daily objective as well as the customers' load model. Section III describes the multi-armed bandit formulation for the electricity pricing problem, presents the modified TS heuristic, and discusses its performance and reliability. Section IV presents simulation results that showcase the efficacy of the approach. The online Appendix [23] contains a table of notation and proofs.

II. PROBLEM FORMULATION

A. The Aggregator's Objective

The aggregator's main goal is to select dispatch signals to manipulate customer demand according to a given optimization objective that varies daily. Specifically, we consider the case where the aggregator broadcasts a dispatch signal $\mathbf{p}_\tau = [p(t)]_{t=1,\dots,T}$ to the population of customers each day (we use $t = 1, \dots, T$ to index time of day and $\tau = 1, \dots, T$ to index days). The set of dispatch signals available for use by the aggregator is denoted as \mathcal{P} . In this paper, without a loss of generality, we will assume that the dispatch signal sent to customers for load shaping purposes is a real-time pricing (RTP) signal.¹

The aggregator's cost function could cover a broad range of goals including (but not limited to) manipulating the population's load to match a target profile, minimizing the distribution grid's electricity cost from the regional retailer, or solving for the dispatch of multiple generators, if a market is operated at the distribution system level.

In this work, on each day τ , we assume the aggregator's cost function is a fixed and known nonlinear function $f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)$ that depends on the load profile $\mathbf{D}_\tau(\mathbf{p}_\tau)$ of the population in response to the daily broadcasted price \mathbf{p}_τ and a random exogenous parameter vector \mathbf{V}_τ .² The population's load profile on day τ , $\mathbf{D}_\tau(\mathbf{p}_\tau)$, is a $T \times 1$ vector with the

¹The reader should note that this choice is not fundamental to the development of the modified learning heuristics we present in this paper. It only allows us to provide a concrete characterization of the response to dispatch signals by mathematically modeling the customers as cost-minimizing agents equipped with home energy management systems in Section II-D.

²We note that the function f need not have a closed form representation and thus can represent the solution of an economic dispatch problem with multiple generators, which can still be handled through our framework. However, without loss of generality and purely for brevity of notation, here we focus on common distribution systems which usually lack two-sided markets, and thus we focus on load profile manipulation for renewable integration and distribution system protection.

t^{th} element corresponding to the population's power demand during time period t . The exogenous and given $T \times 1$ vector \mathbf{V}_τ varies daily and can correspond to a daily target profile reflecting renewable generation forecasts, weather predictions, and grid conditions. We consider the exogenous vectors to be i.i.d. drawn from a distribution defined on a finite sample space \mathcal{V} , with each outcome drawn with a nonzero probability. We would like the reader to note that this assumption is only made for convenience for our theoretical regret performance guarantee in Theorem 1. In a real-world implementation, the daily exogenous parameters could be correlated across days (e.g., due to weather, seasons, weekday/weekend, etc.). However, this correlation does not affect the safety guarantees of our algorithm or its applicability (i.e., it only affects our formal regret results).

The aggregator must ensure that the broadcasted price signals do not result in load profiles that violate distribution system reliability constraints (e.g., nodal voltage, transformer capacities, or line flow limits). As such, if the aggregator had full information about how the population responds to price signals (i.e., full knowledge of $\mathbf{D}_\tau(\mathbf{p}_\tau)$), the aggregator can solve the following optimization problem on day τ to select the optimal price \mathbf{p}_τ^* :

$$\mathbf{p}_\tau^* = \arg \min_{\mathbf{p}_\tau \in \mathcal{P}} f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) \quad (1)$$

$$\text{s.t. } g_j(\mathbf{D}_\tau(\mathbf{p}_\tau)) \leq 0, \quad \forall j = 1, \dots, J \quad (2)$$

where $g_j(\cdot)_{j=1, \dots, J}$ is used to represent the reliability constraints for the distribution system. We note that these general constraints need not be linear for the proposed Thompson sampling approach.

However, as explained in the introduction, knowledge of customers' price response is unavailable to the aggregator. Recall, 1) the aggregator does not want to directly query customers for their price sensitivities, 2) most customers cannot readily characterize their price sensitivities, and 3) customers might not be willing to share this private information. Accordingly, the aggregator needs a method to sequentially choose daily price signals to simultaneously 1) control their daily incurred cost; 2) learn the customers' price response models; and 3) ensure the distribution system constraints are not violated at any time.

B. Distribution System Operational Constraints

As stated previously, there are various operational constraints within a distribution system that should be met in order to ensure adequate service for customers and to prevent grid failures. In the aggregator's daily optimization in Section II-A, the constraints are formulated as general functions $g_j(\cdot)_{j=1, \dots, J}$. Specifically, these general functions represent distribution system parameters (i.e., the nodal voltage $u_\tau(t)$ and power flow through distribution lines $f_\tau(t)$) that should obey the following constraints:

$$u_\tau(t) \geq u^{\min}, \quad \forall t, \tau, \quad (3)$$

$$u_\tau(t) \leq u^{\max}, \quad \forall t, \tau, \quad (4)$$

$$f_\tau(t) \leq S^{\max}, \quad \forall t, \tau, \quad (5)$$

where u^{\min} , u^{\max} , and S^{\max} correspond to the lower voltage limit, upper voltage limit, and power flow limit, respectively, for the population's connection to the distribution grid. We note that $u_\tau(t)$ and $f_\tau(t)$ can be easily derived from the population's load profile $\mathbf{D}_\tau(\mathbf{p}_\tau)$ (See Section IV-B). Now that we have described the aggregator's objective and the distribution system's constraints, we next describe the customers' load model as well as their price response model.

C. Load Flexibility Model

It is hard to approach the problem of learning the response of a population of customers to complex dispatch signals such as RTP as a complete "black box problem", i.e., by just observing the broadcasted price and the load response. There are many reasons for this, including 1) the existence of random or exogenous parameters which lead to variability in the temporal and geographical behavior of electricity demand; 2) the variability of the control objective on a daily basis (e.g., due to randomness in renewable generation outputs, market conditions, or baseload); and 3) the small size of the set of observations that one can gather compared to the high dimensional structure of the load (there are only 365 days in a year, so only 365 sets of prices can be posted). Hence, in this paper, we will be *exploiting the known physical structure* of the problem and making use of our statistical prior knowledge of how the load behaves to lower the problem dimensionality.

Specifically, to lower the dimensionality for the learning problem, we explore the fact that flexible loads only show limited number of "load signatures" (justified due to the automated nature of load response through home energy management systems, the limited types of flexible appliances, and the common electricity usage patterns that emerge from electricity customers as shown in [10], [24]). Let us assume that electric appliances can belong to a finite number of clusters $c \in \mathcal{C}$. For each cluster c , we denote \mathcal{D}_c as the set of feasible daily power consumption schedules that satisfy the energy requirements of the corresponding appliances. Any power consumption schedule, $[d_c(t)]_{t=1, \dots, T} = \mathbf{D}_c \in \mathcal{D}_c$, would satisfy the daily power needs of an appliance in cluster c . For example, consider a cluster that represents plug-in electric vehicles (EVs) that require E_c kWh in the time interval $[t_1, t_2]$ with a maximum charging rate of ρ_c kW. Accordingly, the set \mathcal{D}_c of daily feasible power consumption schedules is given by:

$$\mathcal{D}_c = \left\{ \mathbf{D}_c \mid \sum_{t=t_1}^{t_2} d_c(t) = E_c; 0 \leq d_c(t) \leq \rho_c \right\}. \quad (6)$$

Another specific cluster example is that of electric appliances that are uninterruptible but can perform load shifting (e.g., a dishwasher cannot be interrupted once it is turned on but the start time of the cycle can be shifted). Let $\Pi_c(\cdot)$ denote the load profile of uninterruptible cluster c appliances once they are turned on. For example, $\Pi_c(\cdot)$ could be a rectangular pulse function that outputs the rated power of the appliance, ρ_c for the duration of the appliance's cycle and 0 otherwise. To relay their load flexibility, cluster c users can specify a time interval $[t_{c,1}, t_{c,2}]$ within which the appliance cycle must start (e.g., a user wants the dishwasher to be finished before dinner).

Thus the home energy management system can calculate the best values for the time shift, denoted by t_c , as long as it lies within the interval $[t_{c,1}, t_{c,2}]$. The set \mathcal{D}_c of daily feasible power consumption schedules for appliances in this cluster is given by:

$$\mathcal{D}_c = \{\mathbf{D}_c | d_c(t) = \Pi_c(t - t_c); t_c \in [t_{c,1}, t_{c,2}]\}. \quad (7)$$

For discussion on characterizing the sets for other flexible appliances, including interruptible [25, Sec. 3.2], non-interruptible [25, Sec.3.4], and thermostatically controlled loads [25, Sec. 3.3], we refer the reader to [25].

By adopting this model, the total power consumption flexibility of a population of customers can be characterized as a function of how many appliances belong to each cluster within the population. Let us denote a_c as the number of appliances in cluster c (note that this will vary on a day by day basis as described in the next section). With this notation, we can write the set of feasible daily power consumption profiles for the population, \mathcal{D} :

$$\mathcal{D} = \sum_{c \in \mathcal{C}} a_c \mathcal{D}_c, \quad (8)$$

where the summation and scalar multiplication operations are defined in the sense of Minkowski addition.³

We would like to note that choosing the number of clusters in the model is a control knob that can be tuned by the aggregator as shown in [26]. Using a higher number of potential appliance clusters will increase the accuracy of the load model (i.e., reduce the quantization error in the reproduction of the individual load profiles) and yield better performance in the daily optimization once the true parameters have been sufficiently learned by the aggregator. However, increasing the number of load clusters increases the size of the problem space and increases the randomness in the customers' daily loads thus slowing down the learning rate of the algorithm. The number of clusters will vary depending on the system being analyzed as well as the aggregator's preferences.

D. Price Response Model

In this section, we discuss how the total population responds to dynamic electricity prices given the load flexibility model in (8) and how clustering is used to reduce the dimensionality of the problem. There are two main ways dynamic pricing affects the power consumption: 1) *Automated per cluster response*: Within each load cluster c (i.e., given pre-specified preferences such as EV charging deadlines or AC temperature set points), we assume that the customer chooses the power consumption profile $\mathbf{D}_c \in \mathcal{D}_c$ that minimizes their electricity cost dependent on the daily broadcasted price \mathbf{p}_τ . For appliances in cluster c on day τ , we assume all will choose the same minimum cost power consumption profile:

$$\tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau) = \arg \min_{\mathbf{D}_c \in \mathcal{D}_c} \sum_{t=1}^T p(t) d_c(t). \quad (9)$$

³For two sets A and B defined on a finite dimensional Euclidean space, the Minkowski sum is defined as $A + B = \{\mathbf{a} + \mathbf{b} \mid \mathbf{a} \in A, \mathbf{b} \in B\}$.

We assume that each appliance will always choose the cost minimizing power consumption profile out of the available profile set to combat the fact that the available profile sets \mathcal{D}_c for each cluster can be infinitely large. Thus, we have effectively reduced the dimensionality of the problem as we know *a priori* how each cluster will respond to each price signal (i.e., each cluster will always select its cost minimizing profile). Due to the automated nature of home energy management systems, each cluster selecting its cost minimizing profile is a reasonable assumption once the customers have defined their flexibility preferences, e.g., the desired charge amounts and deadlines for EVs [27], [28]. 2) *Preference Adjustment*: We also consider the fact that customers may respond to price signals by adjusting their preferences. Consider the following example: two customers (Customer-A and Customer-B) live in the same neighborhood but have different sensitivities to electricity prices. If electricity prices are high on a hot summer day, Customer-A might shutdown their air conditioner to avoid a large electricity bill; however, Customer-B prioritizes comfort over cost-savings, and leaves their air conditioner on, no matter the cost. As shown in the previous example, the number of appliances in each cluster, i.e., a_c in (8), also depends on the daily posted price vector \mathbf{p}_τ , and are now denoted as $a_c(\mathbf{p}_\tau)$.

Combining the *automated per cluster response* and *preference adjustment*, we can define the population's load on day τ in response to the posted price \mathbf{p}_τ as follows:

$$\mathbf{D}_\tau^*(\mathbf{p}_\tau) = \sum_{c \in \mathcal{C}} a_c(\mathbf{p}_\tau) \tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau). \quad (10)$$

As stated before, if the aggregator has full knowledge of the customers' price responses, which reduces to having full knowledge of the preference adjustments $a_c(\mathbf{p}_\tau)$, then the aggregator can pick the daily price vector \mathbf{p}_τ^* in order to shape the population's power consumption according to (1). However, as we cannot assume this, we model the $a_c(\mathbf{p}_\tau)$'s as random variables with parameterized distributions, ϕ_c , based on the posted price signal \mathbf{p}_τ and an unknown but constant parameter vector θ^* . Here, θ^* represents the *true model* for the customers' sensitivity to the price signals. This allows for the complex response of the customer population to be represented as a single vector, thus reducing the dimensionality of the problem. We note that while $a_c(\mathbf{p}_\tau)$ may only take integer values in reality, we believe it is justified to relax this integrality constraint and allow it to take continuous values with large enough appliance population size. With this in mind, we would like to highlight three properties of the price response model:

- 1) The preference adjustment models $a_c(\mathbf{p}_\tau)$ are stochastic and their distributions ϕ_c are parameterized by \mathbf{p}_τ and θ^* . This is due to exogenous factors outside of the aggregator's scope that influence customers' power consumption profiles resulting in a level of stochasticity in the responses to prices (i.e., customers will not respond to prices in the same fashion each day).
- 2) The probability distributions of $a_c(\mathbf{p}_\tau)$ (i.e., ϕ_c) are unknown to the aggregator, i.e., the aggregator does not know the true parameter θ^* of the stochastic model.

- 3) The realizations of $a_c(\mathbf{p}_\tau)$ are not directly observable by the aggregator. The aggregator can only monitor the population's total consumption profile \mathbf{D}_τ and cannot observe the decomposed response of each cluster $a_c(\mathbf{p}_\tau)\tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau)$ independently.

Because we have introduced stochasticity to customers' price response models, we appropriately alter the aggregator's optimization problem for selecting the price signal on day τ to account for the distributions ϕ_c :

$$\begin{aligned} \mathbf{p}_\tau^* &= \arg \min_{\mathbf{p}_\tau \in \mathcal{P}} \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)] & (11) \\ \text{s.t.} \quad & \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [g_j(\mathbf{D}_\tau(\mathbf{p}_\tau)) \leq 0] \geq 1 - \mu, \quad \forall j & (12) \end{aligned}$$

where μ is the aggregator's desired reliability metric for the distribution system constraints. In (11), the aggregator now considers minimizing an expected cost and is subject to probabilistic reliability constraints in (12) that depend on the distributions ϕ_c of the preference adjustment models $a_c(\mathbf{p}_\tau)$.

We note that the formulated chance constraints are enforced with respect to uncertainty in the clusters' price sensitivity parameters, not to the exogenous context vector \mathbf{V}_τ . In this work, we assume the daily exogenous vector is fully known each day and does not add uncertainty to the problem. However, uncertainties in the exogenous vector are important to real-world systems such as the power grid and can be accommodated by our approach by adding external noise to these vectors in the same fashion as noise being added to the population's load. This, of course, would further slow down the learning rate of the algorithm due to the added noise reducing the effectiveness of each posterior update.

Clearly, the aggregator needs to learn the underlying parameters of the stochastic models ϕ_c of how customers respond to price signals in order to select price signals for load shaping initiatives (i.e., the aggregator needs to learn θ^*). Our proposed learning approach and pricing strategy for an electricity aggregator is detailed in the next section.

III. REAL-TIME PRICING VIA MULTI-ARMED BANDIT

A. Multi-Armed Bandit Overview

We utilize the multi-armed bandit (MAB) framework to model the iterative decision making procedure of an aggregator implementing a daily load shaping program [29]–[31]. The MAB problem can be described as a decision making problem where an agent has a set of available actions but can only take one action per round. After an action is taken, the agent experiences a cost that is dependent on the action taken. The agent can only learn about the distribution of costs from each action by experimenting. Throughout this iterative procedure, the agent faces the core dilemma: should the agent *exploit* actions that have yielded small costs, or *explore* actions that have not been tested thoroughly? The goal in a MAB problem is to develop a strategy for selecting actions that balance this trade-off and minimize the cumulative cost over a given time span. More thorough explanation and background of the MAB problem can be found in [32].

For the electricity pricing problem, the MAB framework exemplifies the *exploration-exploitation* trade-off dilemma

faced by an aggregator each day. Namely, should the aggregator choose to broadcast untested prices (i.e., *explore*) to learn more information about the customers? Or should the aggregator choose to broadcast well-performing prices (i.e., *exploit*) to manipulate the daily electricity demand?

To evaluate the performance of an algorithm that aims to tackle the exploration-exploitation trade-off, one commonly examines the algorithm's *regret*. Formally, regret is defined as the cumulative difference in cost incurred over \mathcal{T} days between a clairvoyant algorithm (i.e., the optimal strategy that is aware of the customers' price responses) and any proposed algorithm that does not know the customers' price responses:

$$R_{\mathcal{T}} = \sum_{\tau=1}^{\mathcal{T}} f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) - f(\mathbf{D}_\tau(\mathbf{p}^*), \mathbf{V}_\tau). \quad (13)$$

Instead of considering the cumulative difference in objective function value, an alternative metric for regret is to count the number of times that suboptimal price signals are selected over the \mathcal{T} days. For this, we introduce the following notation: let $\mathbf{p}^{\mathbf{V}_\tau, \theta^*}$ denote the optimal price signal for the true model of the population's price response θ^* when the daily exogenous parameter \mathbf{V}_τ is observed on day τ . Any price signal $\mathbf{p}_\tau \neq \mathbf{p}^{\mathbf{V}_\tau, \theta^*}$ is considered a suboptimal price. Moreover, we denote $N_\tau(\mathbf{p}, \mathbf{V})$ as the number of times up to day τ that the algorithm simultaneously observes the exogenous parameter \mathbf{V} and selects the price signal \mathbf{p} . As such, the total number of times that suboptimal price signals are selected over \mathcal{T} days is:

$$\sum_{\mathbf{V} \in \mathcal{V}} \sum_{\mathbf{p} \in \{\mathcal{P} \setminus \mathbf{p}^{\mathbf{V}, \theta^*}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) = \sum_{\tau=1}^{\mathcal{T}} \mathbb{1}\{\mathbf{p}_\tau \neq \mathbf{p}^{\mathbf{V}_\tau, \theta^*}\}, \quad (14)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function that is set equal to one if the criteria is met and zero otherwise. Subsequently, in an iterative decision making problem such as this, the question arises: *how can an aggregator learn to price electricity with bounded regret, and what are the regret bounds we can provide for a proposed algorithm given dynamically changing grid conditions and reliability constraints?* In the following sections, we present a modified Thompson sampling heuristic for the electricity pricing problem to simultaneously learn the true model θ^* for the population, select the daily price signals, ensure grid reliability, and provide a regret guarantee.

B. Thompson Sampling

Thompson sampling (TS) is a well-known MAB heuristic for choosing actions in an iterative decision making problem with the *exploration-exploitation* dilemma [32]–[34]. Two other well-studied frameworks, greedy algorithms and upper-confidence bound (UCB) algorithms, have shown promise in this problem area. However, greedy algorithms are inferior to Thompson sampling in regret performance and UCB algorithms are restricted to simpler linear optimizations [35]–[37], whereas Thompson sampling can readily handle more general objective functions such as those adopted in our paper [38]. Additionally, a novel aspect of our paper is that we have shown how to modify the Thompson sampling heuristic to

account for reliability constraints with a theoretical guarantee (Proposition 1). There are no other bandit optimization approaches known to be able to handle general objective functions with safety constraints. Relevant works here include the analysis of the performance of the UCB algorithm in the linear MAB setting with linear safety-constraints [39], and well as linear TS with linear constraints [40]. In the latter work, it is shown that in the linear case, the presence of linear constraints do not negatively affect the regret performance of TS, which is remarkable and could be a preliminary justification as to why TS performs well in our paper in the presence of general (non-linear) cost and constraint functions.

Simply put, the integral characteristic of Thompson sampling is that the algorithm's knowledge on day τ of the unknown parameter θ^* is represented by the prior distribution $\pi_{\tau-1}$. Each day the algorithm samples $\tilde{\theta}_\tau$ from the prior distribution, and selects an action assuming that the sampled parameter is the true parameter. The algorithm then makes an observation dependent on the chosen action and the hidden parameter and performs a Bayesian update on the parameter's distribution π_τ based on the new observation. Because TS samples parameters from the prior distribution, the algorithm has a chance to explore (i.e., draw new parameters) and can exploit (i.e., draw parameters that are likely to be the true parameter) through out the run of the algorithm.

C. Constrained Thompson Sampling

In this section, we present the MAB heuristic titled Con-TS-RTP adopted to the electricity pricing problem. Con-TS-RTP is a modified Thompson sampling algorithm where the daily optimization problem is subject to constraints (standard TS algorithms do not have constraints)[41].

When initializing π_0 , the initial distribution on the customers' unknown parameter can be selected by the aggregator. If the aggregator has access to prior information regarding the true parameter, then they could initialize the prior as a distribution of their choice. However, if the aggregator has no prior knowledge, a uniform distribution among all available parameters may be used to model the lack of knowledge of the aggregator.

Each day, the algorithm observes the daily target profile \mathbf{V}_τ , draws a parameter $\tilde{\theta}_\tau$ from the prior distribution, broadcasts a price signal to the customers, observes the load profile of the population in response to the broadcasted price, and then performs a Bayesian update on the parameter's distribution π_τ based on the new observation. We note that there are no restrictions on the class of optimization problem to be solved each day; however, in order for our regret guarantee to hold, the aggregator *must* be able to find the globally optimal solution and can use any desired solution method to do so. In our experimental examples, we assume that θ 's and \mathbf{p}_τ 's are chosen from discrete sets in order to be able to guarantee that an enumeration method could solve for the globally optimal price signals each day in spite of non-convexities that arise.

The observation on day τ is denoted as $\mathbf{Y}_\tau = \mathbf{D}_\tau^*(\mathbf{p}_\tau)$ and we assume that each \mathbf{Y}_τ comes from the observation space \mathcal{Y} that is known a priori. When performing the Bayesian update, the algorithm makes use of the following *likelihood function*:

Algorithm 1 CON-TS-RTP

Input: Parameter set Θ ; Price set \mathcal{P} ; Observation set \mathcal{Y} ; Voltage constraints u^{min}, u^{max} ; Power flow constraint S^{max} , Reliability metrics μ, ν

Initialize π_0 based on aggregator's available prior knowledge of customer sensitivity.

- 1: **for** Day index $\tau = 1 \dots \mathcal{T}$ **do**
- 2: Sample the daily hidden parameter $\tilde{\theta}_\tau \in \Theta$ from the aggregator's prior distribution $\pi_{\tau-1}$.
- 3: Observe the daily exogenous parameter \mathbf{V}_τ .
- 4: Broadcast the daily price signal:

$$\hat{\mathbf{p}}_\tau = \arg \min_{\mathcal{P}} \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) | \theta = \tilde{\theta}_\tau]$$

Subject to:

Constraint Set A:

$$\begin{cases} \mathbf{A.1:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{min} | \theta = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \\ \mathbf{A.2:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{max} | \theta = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \\ \mathbf{A.3:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{max} | \theta = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \end{cases}$$

Constraint Set B:

$$\begin{cases} \mathbf{B.1:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{min} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \\ \mathbf{B.2:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{max} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \\ \mathbf{B.3:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{max} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \end{cases}$$

- 5: Observe the population's load response to price \mathbf{p}_τ : $\mathbf{Y}_\tau = \mathbf{D}_\tau^*(\mathbf{p}_\tau)$.
- 6: Update the aggregator's knowledge of the true parameter in the posterior:

$$\forall S \subseteq \Theta: \pi_\tau(S) = \frac{\int_S \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \theta) \pi_{\tau-1}(d\theta)}{\int_\Theta \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \theta) \pi_{\tau-1}(d\theta)}$$

- 7: **end for**
-

$\ell(\mathbf{Y}_\tau; \mathbf{p}, \theta) = \mathbb{P}_\theta(\mathbf{D}_\tau^*(\mathbf{p}_\tau) = \mathbf{Y}_\tau | \mathbf{p}_\tau = \mathbf{p})$. This function calculates the likelihood of observing a specific load profile when broadcasting price \mathbf{p} and the true parameter is θ . The pseudocode for Con-TS-RTP applied to the constrained electricity pricing problem is presented in Algorithm 1.

D. Discussion on Regret Performance of Con-TS-RTP

The regret analysis of Con-TS-RTP is inspired by the results in [38] for TS with nonlinear cost functions. The authors in [42] extended the regret results from [38] by analyzing the effects of an objective function that is dependent on exogenous parameters such as \mathbf{V}_τ . The analysis in the aforementioned papers provides bounds on the total number of times that sub-optimal price signals selected by the algorithm over \mathcal{T} days as specified in equation (14). The regret guarantee we provide in this work extends the result further, allowing for constraints in the daily optimization that are dependent on the sampled $\tilde{\theta}_\tau$. As such, our regret guarantee applies to the Con-TS-RTP algorithm with constraints as formulated in *Constraint Set A* in Algorithm 1. We refer the reader to the online Appendix [23] as well as [42] and [38] for further discussion on the derivation of Theorem 1.

Assumption 1 (Finitely Many Price Signals, Observations): $|\mathcal{P}|, |\mathcal{Y}| < \infty$.

Assumption 2 (Finite Prior, "Grain of truth"): The prior distribution π is supported over finitely many particles: $|\Theta| < \infty$. The true parameter exists within the parameter space: $\theta^* \in \Theta$. The initial distribution π_0 has non-zero mass on the true parameter θ^* (i.e., $\mathbb{P}_{\pi_0}[\theta^*] > 0$).

Assumption 3 (Unique Optimal Price Signal): There is a unique optimal price signal $\mathbf{p}^{\mathbf{V},*}$ for each exogenous parameter $\mathbf{V} \in \mathcal{V}$.

Theorem 1: Under assumptions 1-3 and Constraint Set A in Algorithm 1, for $\delta, \epsilon \in (0, 1)$, there exists $\mathcal{T}^* \geq 0$ s.t. for all $\mathcal{T} \geq \mathcal{T}^*$, with probability $1 - \delta$:

$$\sum_{\mathbf{V} \in \mathcal{V}} \sum_{\mathbf{p} \in \{\mathcal{P} \setminus \mathbf{p}^{\mathbf{V},*}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) \leq B + C(\log \mathcal{T}), \quad (15)$$

where $B \equiv B(\delta, \epsilon, \mathcal{P}, \mathcal{Y}, \Theta)$ is a problem-dependent constant that does not depend on \mathcal{T} , and $C(\log \mathcal{T})$ depends on \mathcal{T} , the sequence of selected price signals, and the Kullback-Leibler divergence properties of the bandit problem (i.e., the marginal Kullback-Leibler divergences of the observation distributions $\text{KL}[\ell(\mathbf{Y}; \mathbf{p}, \theta^*), \ell(\mathbf{Y}; \mathbf{p}, \theta)]$ (The complete description of the $C(\log \mathcal{T})$ term is left to the Appendix).

Proof: The proof is in the online Appendix [23]. ■

In the next section, we discuss the distribution system reliability issue that arises from how the Con-TS-RTP algorithm handles the distribution system constraints (i.e., *Constraint Set A*) and a modification to the Con-TS-RTP algorithm to ensure the constraints are enforced (i.e., *Constraint Set B*).

E. Con-TS-RTP With Improved Reliability Constraints

In order for the aggregator to ensure safe operation of the distribution grid while running the Con-TS-RTP algorithm, the reliability constraints need to hold for the true price response model θ^* each day. However, with the constraints formulated as in Algorithm 1's *Constraint Set A*, the distribution system constraints are only enforced for the sampled $\tilde{\theta}_\tau$ and not necessarily the true parameter θ^* . This entails that the distributions $\{\phi_c\}_{c \in \mathcal{C}}$ are parameterized by the sampled $\tilde{\theta}_\tau$; therefore, they are inaccurate if any parameter $\tilde{\theta}_\tau \neq \theta^*$ is sampled. This could potentially lead to many constraint violations throughout the run of the algorithm resulting in inadequate service for the customers and grid failures.

Due to the importance of reliable operation of the distribution system, we present a modification to the Con-TS-RTP algorithm (i.e., replacing *Constraint Set A* with *Constraint Set B* in Algorithm 1) to increase the reliability of the selected prices and resulting load profiles with respect to the grid constraints. Specifically, we propose alternate constraints that depend on the algorithm's current knowledge of the true parameter, instead of the sampled parameter. In other words, instead of depending on $\tilde{\theta}_\tau$, the proposed alternate constraints depend on the prior distributions $\pi_{\tau-1}$ as follows:

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{\min} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (16)$$

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{\max} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (17)$$

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{\max} | \theta \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (18)$$

where ν is a small constant (detailed in Proposition 1). When considering constraints (16)-(18) in the Con-TS-RTP algorithm, the algorithm will select more conservative price signals each day that can guarantee the distribution system's constraints are met with high probability by using the information in the updated prior distributions. Before analyzing the modified algorithm's reliability, we make the following assumption:

Assumption 4: There exists $\xi^* > 0$ and $\lambda \geq 0$, such that for all $\theta \neq \theta^*$, $\text{KL}[\ell(\mathbf{Y}; \mathbf{p}, \theta^*), \ell(\mathbf{Y}; \mathbf{p}, \theta)] \geq \xi^*$, where

$$\xi_{\theta, \mathbf{p}}^* = \max_{x \in \mathbb{Z}^{>0}} \left\{ \frac{-\lambda}{x} - \frac{4}{\sqrt{x}} \sqrt{\frac{\log |\mathcal{Y}| |\mathcal{P}|}{\delta} + \frac{\log x}{2}} \right. \\ \left. \times \sum_{\mathbf{Y} \in \mathcal{Y}} \left| \log \frac{\ell(\mathbf{Y}; \mathbf{p}, \theta^*)}{\ell(\mathbf{Y}; \mathbf{p}, \theta)} \right| \right\}$$

and

$$\xi^* = \max_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}} \xi_{\theta, \mathbf{p}}^*.$$

Assumption 4 ensures that as the aggregator performs the steps in Algorithm 1, the algorithm's Bayesian updates of the prior distribution π_τ will likely never decrease the mass of the true parameter θ^* below a certain threshold. Specifically, with Assumption 4, it can be shown (as in [38]) that with probability $1 - \delta\sqrt{2}$ the following holds for all $\tau \geq 1$:

$$\pi_\tau(\theta^*) \geq \pi_0(\theta^*) e^{-\lambda |\mathcal{P}|}, \quad (19)$$

where $\lambda \geq 0$ is a chosen parameter (from Assumption 4) that dictates the minimum reachable mass of the true parameter via Bayesian updating. With the modified constraints (16)-(18) and the minimum mass of the true parameter specified in (19), the reliability of Con-TS-RTP can be characterized as follows:

Proposition 1: Under assumptions 1-4, with ν in equations (16)-(18) chosen such that $\nu \leq \mu \pi_0(\theta^*) e^{-\lambda |\mathcal{P}|}$, with probability $1 - \delta\sqrt{2}$, the Con-TS-RTP algorithm with Constraint Set B will uphold the probabilistic distribution system constraints as formulated in (12) for each day $\tau = 1, \dots, \mathcal{T}$.

Proof: The proof is in the online Appendix [23]. ■

Remark: The novelty of Con-TS-RTP is that we can ensure with high probability an unsafe price signal is never selected. We can tune the safety parameter to determine what level of risk is acceptable to the aggregator. We note that the selection of an unsafe price signal has no effect on the learning capability of the algorithm. The Con-TS-RTP algorithm will learn regardless of safe/unsafe price signals. The algorithm will never crash/stop prematurely due to the selection of an unsafe price signal; however, the local distribution grid might surpass safety limits on transformers or line flow limits due to an unsafe price selection on select very limited days, at which points protective measures (e.g., relays) should be used to ensure physical grid safety. We note that this is natural for any learning algorithm dealing with stochastic conditions and unknown system parameters. Contingencies can never be avoided 100%, similar to other grid operation paradigms

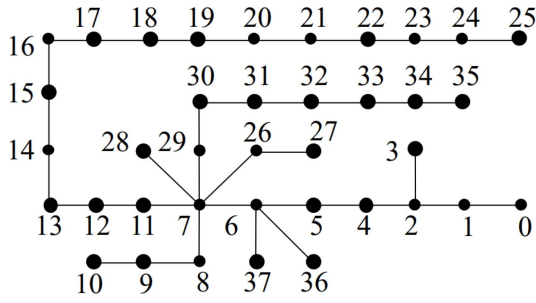


Fig. 1. Radial distribution system.

that deal with uncertain conditions (e.g., wholesale market dispatch with renewables or possible transmission system contingencies). They could only be avoided with a certain high probability when making dispatch decisions. However, it is understood that other protective measures should always be put in place to avoid physical system damage in case of contingencies.

IV. EXPERIMENTAL EVALUATION

A. Test Setup: Radial Distribution System

In this section we describe the power distribution system and the corresponding network parameters for the test case. We consider an actual radial distribution system from the ComEd service territory in Illinois, USA (adopted from [43] and shown in Fig. 1) represented by the undirected graph \mathcal{G} which includes a set of nodes (vertices) \mathcal{N} and a set of power lines (edges) \mathcal{L} . In this work, we consider each node as one population with its own daily load profile; however, each node could be an aggregation of smaller entities downstream of the local distribution connection point. The undirected graph is organized as a tree, with the root node representing the distribution system's substation where it is connected to the regional transmission system. We denote N as the total number of nodes in the network excluding the root node. The nodes are indexed as $i = 0, \dots, N$, and the node corresponding to $i = 0$ (i.e., the root node) is the substation. The power lines are indexed by $i = 1, \dots, N$ where the i -th line is directly upstream of node i (i.e., line i feeds directly to node i). In the following, we denote the parent vertex of node i as \mathcal{A}_i and the set of children vertices of node i as \mathcal{K}_i .

Furthermore, we assume the aggregator has access to measurement data at each node's local connection point. Specifically, the aggregator measures the active power demands at each node i at time t on day τ denoted as $d_{i,\tau}^P(t)$. In order to ensure the delivered power is suitable for the electricity customers, the aggregator also monitors node i 's local voltage at time t on day τ denoted as $v_{i,\tau}(t)$. In the following, we denote the active power daily load profile of node i on day τ as $\mathbf{D}_{i,\tau}^P = [d_{i,\tau}^P(t)]_{t=1,\dots,T}$. Additionally, the aggregator records the active power flows $f_{i,\tau}^P(t)$ on each line $i \in \mathcal{L}$. We note that reactive power should also be monitored in distribution systems, even though it is generally not priced and customers do not consider it in determining their optimal load response to prices. As such, we use the superscript Q for the reactive power at a node, $d_{i,\tau}^Q(t)$, and for reactive power flow

TABLE I
DISTRIBUTION SYSTEM PARAMETERS

| Line | R ($10^{-3}\Omega$) | X ($10^{-3}\Omega$) | S^{max} (KVA) | Line | R ($10^{-3}\Omega$) | X ($10^{-3}\Omega$) | S^{max} (KVA) |
|------|--------------------------|--------------------------|--------------------|------|--------------------------|--------------------------|--------------------|
| 1 | 24.2 | 48.2 | 54 | 20 | 129.5 | 30.9 | 10.8 |
| 2 | 227.3 | 743.5 | 84 | 21 | 15.1 | 5.4 | 14.4 |
| 3 | 76.3 | 18.2 | 10.8 | 22 | 50.8 | 12.1 | 10.8 |
| 4 | 43.6 | 142.7 | 84 | 23 | 69.1 | 16.5 | 10.8 |
| 5 | 25.8 | 84.4 | 84 | 24 | 31.6 | 11.2 | 14.4 |
| 6 | 10.5 | 10.7 | 40.2 | 25 | 96.3 | 23 | 10.8 |
| 7 | 23.2 | 23.6 | 40.2 | 26 | 110.7 | 112.6 | 40.2 |
| 8 | 75.1 | 26.7 | 14.4 | 27 | 2.1 | 0.7 | 14.4 |
| 9 | 114.4 | 27.3 | 10.8 | 28 | 242.1 | 86.2 | 14.4 |
| 10 | 110.8.3 | 67.7 | 14.4 | 29 | 27.3 | 27.8 | 40.2 |
| 11 | 63.7 | 22.7 | 14.4 | 30 | 174.6 | 62.1 | 16.2 |
| 12 | 278.7 | 99.2 | 14.4 | 31 | 43 | 15.3 | 10.8 |
| 13 | 254.2 | 10.8.5 | 14.4 | 32 | 207.8 | 74 | 10.8 |
| 14 | 21.8 | 5.2 | 10.8 | 33 | 109.4 | 38.9 | 14.4 |
| 15 | 57.3 | 20.4 | 14.4 | 34 | 50.5 | 18 | 14.4 |
| 16 | 126.7 | 45.1 | 14.4 | 35 | 165.2 | 58.8 | 14.4 |
| 17 | 48.6 | 11.6 | 10.8 | 36 | 49.5 | 17.6 | 14.4 |
| 18 | 95.1 | 22.7 | 10.8 | 37 | 5.8 | 2.1 | 14.4 |
| 19 | 137.3 | 32.8 | 10.8 | | | | |

on a line, $f_{j,\tau}^Q(t)$. Each line in the distribution system has its own internal resistance denoted as R_i , reactance denoted as X_i , and power limit denoted as S_i^{max} . The parameters for the distribution system are listed in Table I.

B. Power Flow Model

In order to solve for the power flow and nodal voltages of the power distribution system, we make use of the *LinDistFlow* model[44], which is a linear approximation for the AC power flow model.⁴ The *LinDistFlow* model has been extensively studied and verified to be competitive to the nonlinear AC flow model on many realistic feeder topologies including radial [46]–[49]. The *LinDistFlow* model reduces computational complexity by making use of the following linear power flow and voltage equations:

$$d_{i,\tau}^P(t) + \sum_{j \in \mathcal{K}_i} f_{j,\tau}^P(t) = f_{\mathcal{A}_i,\tau}^P(t); \quad \forall t, \tau, i, \quad (20)$$

$$d_{i,\tau}^Q(t) + \sum_{j \in \mathcal{K}_i} f_{j,\tau}^Q(t) = f_{\mathcal{A}_i,\tau}^Q(t); \quad \forall t, \tau, i, \quad (21)$$

$$u_{\mathcal{A}_i,\tau}(t) - 2(f_{i,\tau}^P(t)R_i + f_{i,\tau}^Q(t)X_i) = u_{i,\tau}(t); \quad \forall t, \tau, i \quad (22)$$

where (20) accounts for active power and (21) accounts for reactive power. In (22) we make use of the operator $u_{i,\tau}(t) = (v_{i,\tau}(t))^2$ to provide a linear voltage drop relationship across the distribution system. For the scope of this work, we assume that the substation connection to the regional transmission system (node $i = 0$) is regulated and has a fixed voltage $v_{0,\tau}(t) = 12.5\text{kV}, \forall t, \tau$.

⁴The reader should note that the proposed learning approach is not limited to the *LinDistFlow* model. There are other power flow models that can be utilized such as [45].

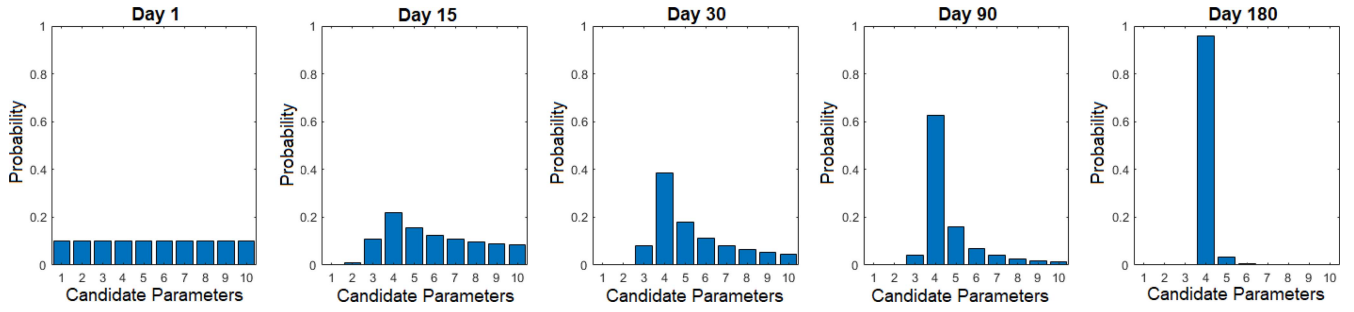


Fig. 2. The 5 plots above portray the evolution of the aggregator's knowledge of the population's hidden parameter at node 10 throughout the learning procedure. The true parameter is parameter 4. From left to right: Day 1 (initialized to uniform distribution, i.e., no knowledge of the true parameter), Day 15 prior, Day 30 prior, Day 90 prior, and Day 180 prior. At day 180, the aggregator is about 95% certain that parameter 4 is the true parameter.

C. Distribution System Operational Constraints

The nodal voltages and line flows calculated in (20)-(22) should obey the following constraints for reliable operation:

$$u_{i,\tau}(t) \geq u_i^{\min}, \quad \forall t, \tau, i \in \mathcal{N}, \quad (23)$$

$$u_{i,\tau}(t) \leq u_i^{\max}, \quad \forall t, \tau, i \in \mathcal{N}, \quad (24)$$

$$f_{i,\tau}^P(t)^2 + f_{i,\tau}^Q(t)^2 \leq (S_i^{\max})^2, \quad \forall t, \tau, i \in \mathcal{L}, \quad (25)$$

where (23)-(24) are the nodal voltage constraints and (25) corresponds to the power constraints for each distribution line.

D. Load Model and Multi-Armed Bandit Formulation

In this test case, we consider 6 time slots each day, each 4 hours long. We consider 10 unique target load profile vectors, with the daily target profile \mathbf{V}_τ for day τ drawn from a uniform distribution each morning. Each of the 10 target load profile vectors corresponds to a desired load curve to accommodate different levels of forecasted renewable generation. Furthermore, the aggregator transmits daily price signals $\mathbf{p}_{i,\tau}$ to each node within the system. The aggregator has a high and low price for each of the 6 time slots resulting in 2^6 possible daily price signals. Since the aggregator is shaping the electricity demand at each node within the distribution system, each node has its own cost $f(\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}), \mathbf{V}_\tau)$ that is dependent on the node's daily demand and the target profile. In this test case, we assume the cost function for each node is the squared deviation of the node's electricity demand from the target profile: $f(\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}), \mathbf{V}_\tau) = |\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}) - \mathbf{V}_\tau|^2$, thus equally penalizing over-usage and under-usage of electricity. We note that the units are KW^2 and if the aggregator had a converting function for the squared deviation (KW^2) to $\text{\$U.S.D.}$, then we could calculate the monetary losses of the system. In our experimental examples, we make use of discrete sets for the available θ 's and \mathbf{p}_τ 's to guarantee that an enumeration-based method could solve for the globally optimal price signals each day in spite of problem non-convexities.

We consider 20 unique load flexibility clusters in this test case. Each cluster's parameters represent the varying start/stop times, total energy demands, and power limitations common to EV loads in residential areas and are of the form presented in equation (6). We note that we generated the population's load price response directly using the same clustering model (i.e., the actual load response in the simulation is at the level

of 20 clusters and can be well represented by the 20 clusters *plus additive noise*. For a discussion on the effects of poor clustering, we refer the reader to Section IV-F). Each node in the distribution system is comprised of these 20 load clusters with its own unique sensitivities $a_{i,c}(\mathbf{p}_\tau)$ for each cluster. Each sensitivity parameter is selected as $a_{i,c}(\mathbf{p}_{i,\tau}) \sim \mathcal{N}(\frac{\beta_c}{\theta_i^* \mathbf{p}_{i,\tau}}, \sigma^2)$ each day where β_c is a cluster specific constant known by the aggregator (we note that β_c represents *a priori* knowledge of customers' preferences and could come from behavioral studies; however, our framework does not require this and β_c can be completely omitted in cases where prior information is unavailable). Each node's price sensitivity, i.e., parameter to be learned, θ_i^* , is a vector of length 6 and the set of possible parameters, Θ , contains 10 unique vectors. Unless noted, the reliability parameter chosen for the Con-TS-RTP algorithm is $\nu = 0.1$.

Note on reactive power: We note that reactive power is generally price insensitive; however, reactive power is present in a distribution system and affects the constraints of the system. Reactive power flows alter how the price sensitive loads are limited by the operational constraints of the system (i.e., active and reactive flows on lines affect the capacity available for the price responsive loads). Due to the lack of data as to how much reactive power is present in the distribution system due to our appliance clusters and otherwise, for our numerical examples, we omit the inclusion of reactive power to only view the appliance clusters' active load profiles within the distribution system. For further discussion on this, we refer the reader to papers that fully capture the effects of reactive power in such problems such as [50] and [51] in which the authors showcase techniques to handle distribution systems with chance constraints.

In the following sections, without loss of generality, we assume that reactive power is not responsive to the pricing signals. We note that our proposed learning approach can accommodate reactive power flows (*LinDistFlow* can as well); however, our goal was to show proof of concept of our learning/pricing approach with active customer loads, thus reactive power flow will be examined in future work.

E. Results

We simulated the Con-TS-RTP algorithm for 365 days for an aggregator attempting to learn the sensitivities of the nodes

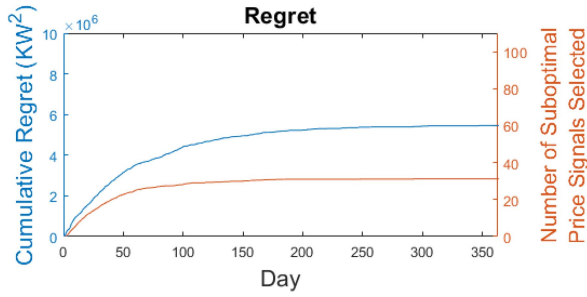


Fig. 3. Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that the y-axis (left) units are KW^2 for the squared load deviation from the target profile.

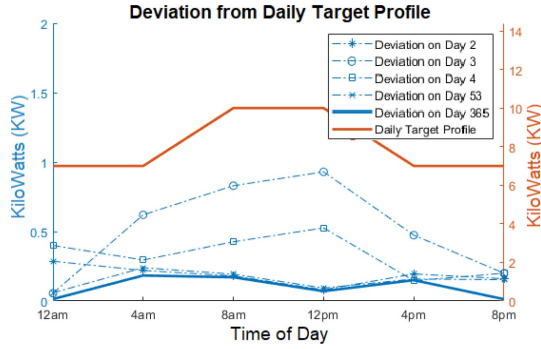


Fig. 4. Deviation of node 10's demand from a specific daily target profile.

in the system and shape their demands. In the following, we highlight the results of the simulation at node 10 of the radial distribution system. Figure 2 presents the evolution of the prior distribution for node 10's hidden parameter. On day 1, the prior was initialized to a uniform distribution among the candidate parameters, and by day 180 the weight on the true parameter exceeded 0.95.

Figure 3 presents the regret performance of Con-TS-RTP at node 10. As seen in Figure 3, the regret curve flattens after day 130 as the algorithm never chooses a suboptimal price signal after this day.

Figure 4 presents node 10's deviation from a specific daily target profile. On days 2, 3, 4, 53, and 365 the same target profile (i.e., $\mathbf{V}_2 = \mathbf{V}_3 = \mathbf{V}_4 = \mathbf{V}_{53} = \mathbf{V}_{365}$) was drawn and the aggregator selected different price signals to shape the node's demand. As seen in Fig. 4, the deviation from the target profile on day 365 is less than the deviation on the other days as the algorithm has learned the true parameter and selects the optimal price signal to shape the load.

In Figure 5, we present the distribution system constraint violations that were avoided by using Con-TS-RTP instead of an unconstrained TS algorithm. Clearly, in the early learning stages, the unconstrained TS algorithm does not have accurate knowledge of the hidden parameters and violates the distribution system constraints often. Con-TS-RTP is more conservative with its exploration of untested price signals and avoids the constraint violations made by the unconstrained TS algorithm. Last, we note that the simulation was implemented in MATLAB and CVX on an i7 processor with 16gb of RAM. The 365 day simulations were run in less than 5 minutes.

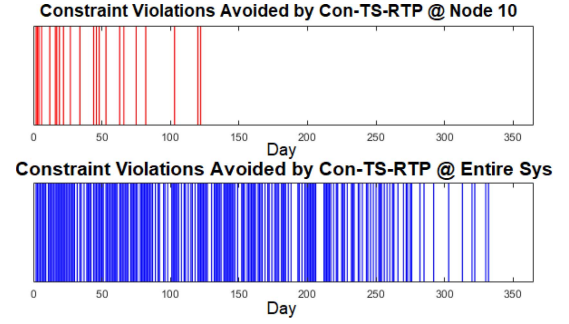


Fig. 5. Top: Distribution system constraint violations at node 10 avoided by using Con-TS-RTP instead of an unconstrained TS. Bottom: Distribution system constraint violations across the entire system avoided by using Con-TS-RTP instead of an unconstrained TS.

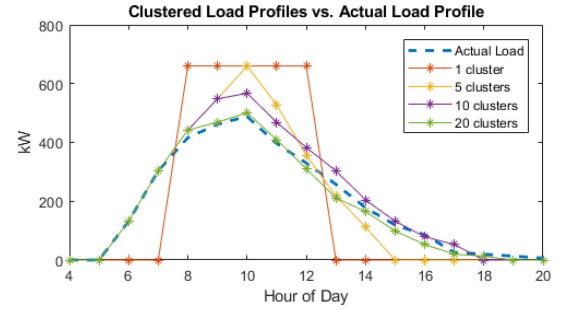


Fig. 6. Effects of changing the number of clusters to model an actual load. Specifically, load profiles for 4 cluster models compared to the actual load profile for a population of 100 charging EVs.

F. Effects of Clustering

In this section, we portray the effects of selecting different numbers of clusters to represent a true load as well as the effects of selecting too few clusters on the performance of our Con-TS-RTP algorithm. First, in Figure 6, we perform a simple demonstration. We considered a population of 100 EVs with random charging requests and then constructed clusters to view the accuracy of the clustered load profiles versus the actual load profile. As shown in Figure 6, using 1, 5, or 10 clusters to represent the EV population results in load profiles quite different from the actual; however, with 20 clusters, the load profile begins to match the actual profile.

Furthermore, in Figure 7 we show the effects of reducing the number of clusters in the load model on the regret performance of our Con-TS-RTP algorithm. Specifically, we focus on the same setup as Section IV-D with the exception that we have the Con-TS-RTP algorithm use a 10 cluster model instead of the 20 cluster model for the population to see the effects of an inaccurate cluster model. As shown in Figure 7, the regret curve for this case never flattens and the algorithm is never able to select the optimal price signal. This is because the algorithm's model of the load (i.e., the 10 clusters) is unable to accurately model the population's response and causes the algorithm to select incorrect prices every day.

G. Evolving Price Sensitivity

In this section, we show an example of what happens when customers' sensitivities change over time and how a Bayesian learning approach can naturally adapt and account for these

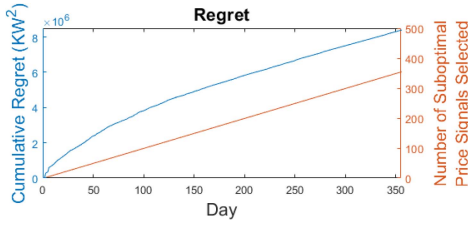


Fig. 7. Effects of using too few clusters for the population’s load model. We show the regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$ for a 10 cluster model instead of the 20 cluster model as previously shown in Fig. 3. Due to the inaccuracies of the 10 cluster model, the algorithm is never able to select the optimal price signals.

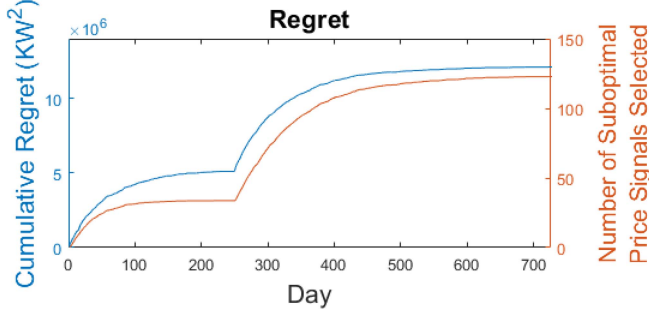


Fig. 8. Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that on day 250, the hidden parameter was altered.

dynamic changes. Specifically, we simulated the same system as in Figure 3, but on day 250, we altered the true θ_i^* parameter. As seen in Figure 8, the regret curves first flatten around day 125, then increase at day 250, and then flatten again near day 325. This shows that Con-TS-RTP was able to successfully learn the first and second true parameter without any modifications to the algorithm. The algorithm naturally shifts its belief about the true parameter as it observes outputs that do not match its current belief.

H. Non-Repeating Target Profiles

In the previous case study, we assumed a low number of target profiles (10 profiles) to satisfy the assumptions we have made for our theoretical results. In this section, we demonstrate how extending the number of target profiles to 365 does not negatively affect the performance of the algorithm in practice. Furthermore, we ensure that once a target profile has been viewed by the aggregator, it is never drawn again. Thus, each day the aggregator is posting a price to shape the population’s load to match a target profile that it has never seen before. As shown in Figure 9, enlarging the set of target profiles does not slow down the learning process. Note that in Figure 9 the regret flattens near trial 100 which matches the duration of the learning period seen in Figure 3 (i.e., in simulation, the aggregator is still able to learn the true parameter when the number of target profiles is increased from 10 to 365, resulting in similar regret curves).

I. Effects of Varying the System Reliability Metric

In this section, we discuss the effects of varying the system reliability parameters in the daily optimization’s constraints (i.e., altering the value of ν for the system constraints

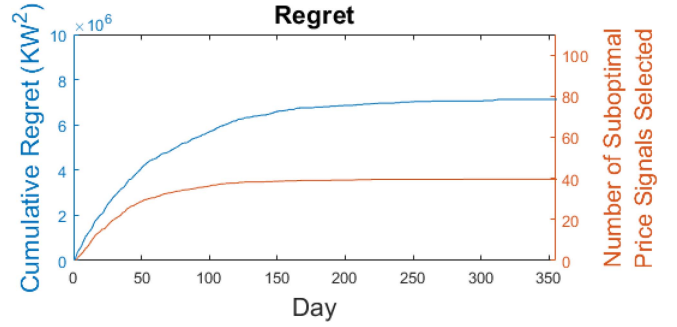


Fig. 9. Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that on each day, the sampled V_t has never been seen by the aggregator.

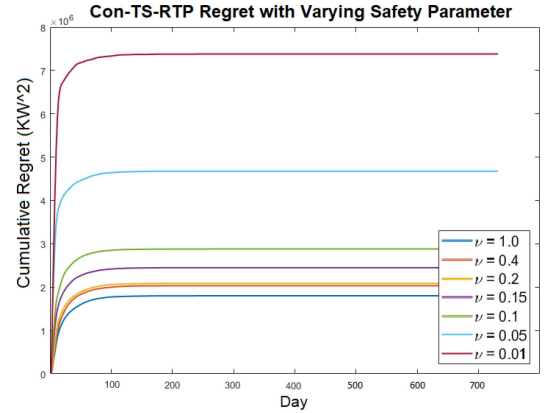


Fig. 10. Regret curves for various system reliability metrics. Each curve is an average of 20 independent simulations.

formulated as in (16)-(18)). As described in Sections II-D and III-E, the reliability metric dictates the aggregator’s allowable probability of a constraint violation under its current belief distribution about the unknown parameter. Decreasing ν is restricting the algorithm to avoid violations and setting $\nu = 1$ is equivalent to solving the daily optimization without the constraints altogether. In Figure 10, we simulated the system with varying reliability parameters. Specifically, each curve shown is the average regret at node 10 over 20 independent simulations. As shown in Figure 10, the regret increases as the desired reliability increases (smaller ν). This is because the aggregator is forced to select more conservative prices during the learning procedure to ensure that the constraints are met with higher probability.

J. Comparison With Two-Stage Learning

In this section, we present a comparison of the Con-TS-RTP approach versus a 2-stage “learn” and then “optimize” algorithm, where the first stage consists of pure exploration and the second stage purely exploits the knowledge gained in the first stage. The simulation setup is the same as the setup used in Section IV-D. A description of the 2-stage algorithm used is as follows: The aggregator decides the duration of the learning stage *a priori*, (in Figure 11, we present regret curves for learning stages with durations of 5, 15, and 25 days) and during this learning stage, the aggregator randomly selects price signals from a predetermined safe set of prices (i.e., prices high

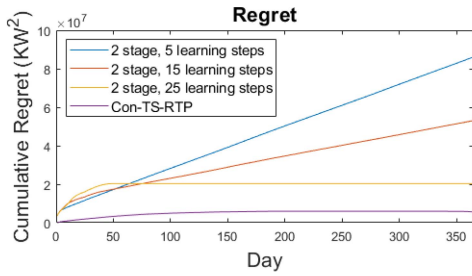


Fig. 11. Regret performance of Con-TS-RTP and a 2-Stage algorithm at node 10 with $\nu = 0.1$. Note that the 5 day (blue) and 15 day (red) learning algorithms were unable to converge to the optimal price signals. The blue and red curves never flatten because their learning stages were too brief to adequately learn the customers' preferences and are unable to select the optimal price signals, resulting in a linearly growing regret. However, the 25 day (yellow) learning stage algorithm is able to adequately learn the population's parameters and select optimal prices after that.

enough such that constraints cannot be violated), observes the populations' responses, and performs posterior updates. Then, after the learning stage is complete, for the remainder of time the aggregator broadcasts the best price signals with respect to the knowledge of the unknown parameter at the end of the learning stage (the selected price signal will ensure safety but might be potentially suboptimal depending on the duration of the learning stage). Clearly, the two most significant shortcomings of the 2-stage approach are: 1) arbitrarily bad performance during the learning stage due to random price selection; and 2) difficulty selecting a sufficient duration of the learning stage. As seen in Figure 11, this 2-stage myopic algorithm results in linear regret in the 5 day and 15 day learning stage curves. Due to an insufficient number of posterior updates, the aggregator is forced to post suboptimal price signals to ensure safety given its noisy knowledge of the unknown parameter after the learning stage is over. On the other hand, the 25 day learning stage is able to converge to the optimal price signals, but the performance during the learning stage causes fast growth of regret whereas Con-TS-RTP is able to avoid all of the aforementioned shortcomings.

V. CONCLUSION

In this paper, we presented a multi-armed bandit problem formulation for an electricity aggregator attempting to run a real-time pricing program for load shaping (e.g., to reduce demand at peak hours, integrate more intermittent renewables, track a desired daily load profile, etc). We made use of a constrained Thompson sampling heuristic, Con-TS-RTP, as a solution to the *exploration/exploitation* problem of an aggregator *passively* learning customers' price sensitivities while broadcasting price signals that influence customers to alter their demand to match a desired load profile. The proposed Con-TS-RTP algorithm permits day-varying target load profiles (i.e., multiple target load profiles reflecting renewable forecasts and desired demand patterns) and takes into account the actual operational constraints of a distribution system to ensure that the customers receive adequate service and to avoid potential grid failures. Additionally, our setup accounts for complex electricity usage patterns of the customers by classifying different load clusters based on electricity demand

and load flexibility. We discussed a regret guarantee for the proposed Con-TS-RTP algorithm which bounds the total number of suboptimal price signals broadcasted by the aggregator. Furthermore, we discussed an operational reliability guarantee that ensures the power distribution system constraints are upheld with high probability throughout the run of the Con-TS-RTP algorithm.

REFERENCES

- [1] C. Eid, E. Koliou, M. Valles, J. Reneses, and R. Hakvoort, "Time-based pricing and electricity demand response: Existing barriers and next steps," *Utilities Policy*, vol. 40, pp. 15–25, Jun. 2016.
- [2] V. Gomez, M. Chertkov, S. Backhaus, and H. J. Kappen, "Learning price-elasticity of smart consumers in power distribution systems," in *Proc. IEEE 3rd Int. Conf. Smart Grid Commun. (SmartGridComm)*, Tainan, Taiwan, 2012, pp. 647–652.
- [3] Z. Xu, T. Deng, Z. Hu, Y. Song, and J. Wang, "Data-driven pricing strategy for demand-side resource aggregators," *IEEE Trans. Smart Grid*, vol. 9, no. 1, pp. 57–66, Jan. 2018.
- [4] P. Li and B. Zhang, "Linear estimation of treatment effects in demand response: An experimental design approach," 2017. [Online]. Available: arXiv:1706.09835.
- [5] P. Li, H. Wang, and B. Zhang, "A distributed online pricing strategy for demand response programs," *IEEE Trans. Smart Grid*, vol. 10, no. 1, pp. 350–360, Jan. 2019.
- [6] K. Khezeli, W. Lin, and E. Bitar, "Learning to buy (and sell) demand response," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6761–6767, 2017.
- [7] L. Jia, L. Tong, and Q. Zhao, "An online learning approach to dynamic pricing for demand response," 2014. [Online]. Available: arXiv:1404.1325.
- [8] K. Khezeli and E. Bitar, "Risk-sensitive learning and pricing for demand response," *IEEE Trans. Smart Grid*, vol. 9, no. 6, pp. 6000–6007, Nov. 2018.
- [9] L. Jia, Q. Zhao, and L. Tong, "Retail pricing for stochastic demand with unknown parameters: An online machine learning approach," in *Proc. 51st IEEE Annu. Allerton Conf. Commun. Control Comput. (Allerton)*, Monticello, IL, USA, 2013, pp. 1353–1358.
- [10] D. Zhou, M. Balandat, and C. Tomlin, "Residential demand response targeting using machine learning with observational data," in *Proc. IEEE 55th Conf. Decis. Control (CDC)*, Las Vegas, NV, USA, 2016, pp. 6663–6668.
- [11] Y. Li, Q. Hu, and N. Li, "Learning and selecting the right customers for reliability: A multi-armed bandit approach," in *Proc. IEEE Conf. Decis. Control (CDC)*, Miami Beach, FL, USA, 2018, pp. 4869–4874.
- [12] D. Kalathil and R. Rajagopal, "Online learning for demand response," in *Proc. IEEE 53rd Annu. Allerton Conf. Commun. Control Comput. (Allerton)*, Monticello, IL, USA, 2015, pp. 218–222.
- [13] E. Dall'Anese, K. Baker, and T. Summers, "Chance-constrained AC optimal power flow for distribution systems with renewables," *IEEE Trans. Power Syst.*, vol. 32, no. 5, pp. 3427–3438, Sep. 2017.
- [14] R. Mieth and Y. Dvorkin, "Data-driven distributionally robust optimal power flow for distribution systems," *IEEE Control Syst. Lett.*, vol. 2, no. 3, pp. 363–368, Jul. 2018.
- [15] R. Mieth and Y. Dvorkin, "Online learning for network constrained demand response pricing in distribution systems," 2018. [Online]. Available: arXiv:1811.09384.
- [16] Y. Wang, Q. Chen, C. Kang, M. Zhang, K. Wang, and Y. Zhao, "Load profiling and its application to demand response: A review," *Tsinghua Sci. Technol.*, vol. 20, no. 2, pp. 117–129, Apr. 2015.
- [17] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Trans. Power Syst.*, vol. 28, no. 1, pp. 430–440, Feb. 2013.
- [18] J. M. Foster and M. C. Caramanis, "Optimal power market participation of plug-in electric vehicles pooled by distribution feeder," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2065–2076, Aug. 2013.
- [19] S. Lin, F. Li, E. Tian, Y. Fu, and D. Li, "Clustering load profiles for demand response applications," *IEEE Trans. Smart Grid*, vol. 10, no. 2, pp. 1599–1607, Mar. 2019.
- [20] W. Zhang, J. Lian, C. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 4655–4664, Nov. 2013.

- [21] C. Chelmiss, J. Kolte, and V. K. Prasanna, "Big data analytics for demand response: Clustering over space and time," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Santa Clara, CA, USA, Oct. 2015, pp. 2223–2232.
- [22] Y. Zhang, W. Chen, R. Xu, and J. Black, "A cluster-based method for calculating baselines for residential loads," *IEEE Trans. Smart Grid*, vol. 7, no. 5, pp. 2368–2377, Sep. 2016.
- [23] N. Tucker, A. Moradipari, and M. Alizadeh, "Constrained thompson sampling for real-time electricity pricing with grid reliability constraints," 2020. [Online]. Available: arXiv:1908.07964.
- [24] J. Kwac, J. Flora, and R. Rajagopal, "Household energy consumption segmentation using hourly data," *IEEE Trans. Smart Grid*, vol. 5, no. 1, pp. 420–430, Jan. 2014.
- [25] M. Alizadeh, A. Scaglione, A. Applebaum, G. Kesidis, and K. Levitt, "Reduced-order load models for large populations of flexible appliances," *IEEE Trans. Power Syst.*, vol. 30, no. 4, pp. 1758–1774, Jul. 2015.
- [26] M. Alizadeh, A. Scaglione, R. J. Thomas, and D. Callaway, "Information infrastructure for cellular load management in green power delivery systems," in *Proc. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, Brussels, Belgium, 2011, pp. 13–18.
- [27] T.-H. Chang, M. Alizadeh, and A. Scaglione, "Coordinated home energy management for real-time power balancing," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, San Diego, CA, USA, 2012, pp. 1–8.
- [28] M. Alizadeh and A. Scaglione, "Least laxity first scheduling of thermostatically controlled loads for regulation services," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Austin, TX, USA, 2013, pp. 503–506.
- [29] A. Krishnamurthy, Z. S. Wu, and V. Syrgkanis, "Semiparametric contextual bandits," 2018. [Online]. Available: arXiv:1803.04204.
- [30] D. J. Foster, A. Agarwal, M. Dudík, H. Luo, and R. E. Schapire, "Practical contextual bandits with regression oracles," 2018. [Online]. Available: arXiv:1803.01088.
- [31] T. Xu, Y. Yu, J. Turner, and A. Regan, "Thompson sampling in dynamic systems for contextual bandit problems," 2013. [Online]. Available: arXiv:1310.5008.
- [32] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, "A tutorial on Thompson sampling," *Found. Trends Mach. Learn.*, vol. 11, no. 1, pp. 1–96, 2018.
- [33] D. Russo and B. Van Roy, "Learning to optimize via posterior sampling," *Math. Oper. Res.*, vol. 39, no. 4, pp. 1221–1243, 2014.
- [34] S. Agrawal and N. Goyal, "Analysis of Thompson sampling for the multi-armed bandit problem," in *Proc. 25th Annu. Conf. Learn. Theory*, 2012, pp. 1–26.
- [35] P. Auer, "Using confidence bounds for exploitation-exploration trade-offs," *J. Mach. Learn. Res.*, vol. 3, pp. 397–422, Mar. 2003.
- [36] V. Dani, T. P. Hayes, and S. M. Kakade, "Stochastic linear optimization under bandit feedback," in *Proc. Annu. Conf. Learn. Theory*, 2008, pp. 355–366.
- [37] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, "Improved algorithms for linear stochastic bandits," in *Advances in Neural Information Processing Systems*. Red Hook, NY, USA: Curran Assoc., Inc., 2011, pp. 2312–2320.
- [38] A. Gopalan, S. Mannor, and Y. Mansour, "Thompson sampling for complex online problems," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 100–108.
- [39] S. Amani, M. Alizadeh, and C. Thrampoulidis, "Linear stochastic bandits under safety constraints," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, 2019, pp. 9252–9262.
- [40] A. Moradipari, S. Amani, M. Alizadeh, and C. Thrampoulidis, "Safe linear thompson sampling with Side Information," 2019. [Online]. Available: arXiv:1911.02156.
- [41] V. Saxena, J. Jaldén, J. E. Gonzalez, I. Stoica, and H. Tullberg, "Constrained thompson sampling for wireless link optimization," 2019. [Online]. Available: arXiv:1902.11102.
- [42] A. Moradipari, C. Silva, and M. Alizadeh, "Learning to dynamically price electricity demand based on multi-armed bandits," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Anaheim, CA, USA, Nov. 2018, pp. 917–921.
- [43] P. Andrianeisis, M. Caramanis, R. Masiello, R. Tabors, and S. Bahramirad, "Locational marginal value of distributed energy resources as non-wires alternatives," *IEEE Trans. Smart Grid*, vol. 11, no. 1, pp. 270–280, Jan. 2020.
- [44] M. E. Baran and F. F. Wu, "Optimal capacitor placement on radial distribution systems," *IEEE Trans. Power Del.*, vol. 4, no. 1, pp. 725–734, Jan. 1989.
- [45] M. D. Sankur, R. Dobbe, E. Stewart, D. S. Callaway, and D. B. Arnold, "A linearized power flow model for optimization in unbalanced distribution systems," 2016. [Online]. Available: arXiv:1606.04492.
- [46] H. J. Liu, "Decentralized optimization approach for power distribution network and microgrid controls," Ph.D. dissertation, Dept. Elect. Comput. Eng., Univ. Illinois Urbana-Champaign, Urbana, IL, USA, 2017.
- [47] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3794–3803, Sep. 2016.
- [48] P. Šulc, S. Backhaus, and M. Chertkov, "Optimal distributed control of reactive power via the alternating direction method of multipliers," *IEEE Trans. Energy Convers.*, vol. 29, no. 4, pp. 968–977, Dec. 2014.
- [49] M. Farivar, L. Chen, and S. Low, "Equilibrium and dynamics of local voltage control in distribution systems," in *Proc. 52nd IEEE Conf. Decis. Control*, Florence, Italy, 2013, pp. 4329–4334.
- [50] M. Lubin, Y. Dvorkin, and L. Roald, "Chance constraints for improving the security of AC optimal power flow," *IEEE Trans. Power Syst.*, vol. 34, no. 3, pp. 1908–1917, May 2019.
- [51] R. Mieth and Y. Dvorkin, "Distribution electricity pricing under uncertainty," *IEEE Trans. Power Syst.*, vol. 35, no. 3, pp. 2325–2338, May 2020.